

Research and Development in the Jayenne Implicit Monte Carlo Project

Todd J. Urbatsch, Allan B. Wollaber,
Kelly G. Thompson, Jeffery D. Densmore,
Gabriel M. Rockefeller, CCS-2; Timothy M. Kelley,
CCS-7

The Jayenne Project, which began in 1997, provides simulation capability for thermal radiative transfer in the X-ray regime for high-energy-density physics applications such as supernova explosions, inertial confinement fusion, and radiation flow experiments at facilities such as SNL's Z-Pinch, the Omega Facility, and the National Ignition Facility [1]. The Jayenne Project uses the Fleck and Cummings [2] Implicit Monte Carlo (IMC) method to simulate the transport of thermal radiation that is non-linearly and tightly coupled to hot material. The Jayenne Project's software is multi-dimensional, runs on Adaptive Mesh Refinement (AMR) meshes, has different parallel schemes, and provides continuously improving production-level simulation while serving as a vehicle for methods research. Some of the recent research items include more accurate sub-cell representation of emission locations, adaptive implicitness, efficiency improvements (both methods- and architecture-based), and variance reduction.


In Implicit Monte Carlo (IMC), particle emission sites need to be chosen within a spatial cell. A piecewise constant representation of the emission energy in each cell produces waves that move too fast. The Jayenne Project software currently utilizes a linear discontinuous tilt scheme based on face averages, which has worked sufficiently well but can produce inaccuracies when cells are too thick and when smaller time steps increase the fraction of volume emission particles. Longer-term research to guide the tilt development has involved analytic moment analysis for continuous transport problems on spatial grids [3]. Recent mid-term research has looked at a sub-cell wavefront representation, which has evolved to a two-segment linear representation with a variable inflection point (VIP) to represent the emission function more accurately everywhere in the problem, not just the wavefront [4]. Figure 1 shows the VIP scheme in 1D, and Fig. 2 shows the improvement in the wavespeed for a spatially under-resolved Marshak Wave.

IMC's results have associated noise, so variance reduction research is an important and recurring activity. One variance reduction technique this year was in the momentum and Eddington tensor tallies [5], which was stimulated by the Jayenne Project's use as the high-order solver in LANL's high-order/low-order research effort [6]. The Jayenne Project

represents curvilinear geometries with Cartesian geometries, such as RZ geometry with a 3D flat-top wedge; the momentum is tallied in Cartesian space and then ensemble-combined to get the radial momentum. By eliminating unnecessary sampling and recasting the estimators as embedded curvilinear estimators, we drastically reduced the variance in each of these tallies. Another variance reduction effort involved material (ion) sources in conjunction with an existing variance/work-reduction capability that refrains from sampling particles in cells whose temperature is below a cutoff. In the IMC method, material sources are sourced partially into the radiation as part of the time-implicitness. Normally, the user can apply a cutoff so that particles are not wasted in large, cold, unimportant regions. If there is any sort of source in any spatial cell, though, we currently disable the cutoff to help conserve energy. We found that, in certain cases, doing this actually detrimentally increased the variance. We have implemented a new capability such that, in any cell below the cutoff that will not go unstable, we bypass the IMC time-implicit treatment and deposit the material source into the material. As the cell heats up, we transition back to the normal IMC material/radiation splits. This new approach reduces variance without sacrificing the stability of the IMC method.

IMC can be slow, and there are several existing algorithms that speed up the IMC. We have been developing a Discrete Diffusion Monte Carlo (DDMC) method that couples with the IMC and more cheaply moves a particle according to a diffusion approximation in regimes where the

Fig. 1. Schematics of the two-segment, variable-inflection-point representation of the emission distribution.






Fig. 2. Marshak wave profiles at $t=10$ sh using $\Delta t=0.001$ sh for the current (top) and variable inflection point (bottom) tilt schemes with varying spatial grids.

diffusion approximation is valid. The DDMC method has been shown to be up to a few orders of magnitude faster than the current state of the art. This year, the DDMC method was extended to a (frequency) multigroup model where particles whose frequencies are below a “thick” cutoff frequency are transported according to DDMC [7].

The Jayenne Project was one of the first software projects running on Roadrunner, and its adaptation focused on RZ geometry. In the summer of 2011, the 3D capability was adapted to Roadrunner with the goal of getting the most out of Roadrunner’s last year of production. Looking beyond Roadrunner, we have begun exploring IMC on graphics processing units (GPU). A first project was to investigate accelerating our source calculation, which involves several loops over the spatial cells. The source calculation is a serial operation in spatial-domain replication parallelism and, in some problems, can quickly dominate runtime with increasing numbers of processes. The most important outcome from this exercise was that a new software component, GPU Device, was added to our underlying component library. This component wraps both CUDA (NVIDIA’s parallel computing architecture) runtime and

driver application programming interfaces (API) to provide an object-oriented view of GPU hardware and GPU kernels. The modularity of this component allows unit testing of low-level GPU capabilities isolated from other portions of the Jayenne Project codes. The new GPU component was used in creating a GPU-optimized version of the Jayenne Project’s Source_Builder algorithm. During this exploration of heterogeneous programming techniques, streaming SIMD extension (SSE) vectorization by hand and OpenMP threads under message-passing interface (MPI) were also explored. For the source builder prototype, the SSE vectorization showed the best performance improvement, but the cost of moving the data to and from the GPU

still overwhelmed the speedup. All of these approaches, each of which requires different coding of the same algorithms, fit nicely into our CMake build system, requiring only subdirectories of the accelerator-specific coding. Moving beyond simple loops, we are using mini-applications to investigate the best design for rewriting our actual particle transport on GPUs.

- [1] Urbatsch, T.J. and T.M. Evans, “Milagro Version 2: An IMC Code for Thermal Radiative Transfer: Capabilities, Development, and Usage,” LANL Technical Report, LA-14195-MS, available at <http://www.osti.gov/bridge> (2005).
- [2] Fleck, J.A. Jr. and J.D. Cummings, *J Comput Phys* **8**, 313 (1971).
- [3] Densmore, J.D., “Spatial Moments of Continuous Transport Problems Computed on Grids,” *Trans. Theory Stat. Phys.* accepted (2011).
- [4] Urbatsch, T.J. et al., “Linear sub-cell representations of emission in particle-based thermal radiation transport,” LANL Technical Memorandum CCS-2:2011-57(U), LA-UR-11-06881 (2011).
- [5] Wollaber, A., “Implementing Variance Reduction in IMC Momentum Tallies,” LANL Technical Memorandum CCS-2:11-10(U) (2011).
- [6] Park, H. et al., “A Consistent, Moment-Based Multiscale Solution Approach for Thermal Radiative Transfer Problems,” *Trans. Theory Stat. Phys.* submitted (2011).
- [7] Densmore, J.D. et al., “A Hybrid Transport-diffusion Monte Carlo Method for Frequency-dependent Radiative-Transfer Simulations,” *J Comput Phys*, to be submitted (2012).

Funding Acknowledgement

DOE NNSA, Advanced Simulation and Computing Program; Weapons Program, Science Campaign 4.6; LANL Laboratory Directed Research and Development Program